dng_validate
Version 1.3
22-Jun-09

"dng_validate" is a command-line tool that parses the tag structure of DNG (and other TIFF-EP based format) files, and reports any deviations from the DNG specification that it finds.

The usage syntax is:

dng_validate [-v] [-d <number>] [-f] [-b4] [-s < CFA index >] [-q <target-binned-width>] [-cs1|-cs2|-cs3|-cs4|-cs5|-cs6] [-16] { [-1 <stage1-out-filename>] [-2 <stage2-out-filename>] [-3 <stage3-out-filename>] [-tif <TIFF-out-filename>] [-dng <DNG-out-file>] <list of files>}*

Any deviations from the DNG specification are written to the standard error stream.

The "-v" option turns on "verbose" mode, which writes the parsed tag structure to the standard output stream. Any tags that are not parsed by this tool are preceded by an asterisk.

The "-d <number>" option both implies verbose mode, and also specifies the maximum number of lines of data displayed per tag.

The "-f" option switches dng_validate to using floating-point math where possible, instead of the default 16-bit integer.

The "-b4" option causes the demosaic algorithm to produce a four channel output rather than a three channel one. (The input DNG must be a three channel Bayer pattern image.) This option is only useful when used with the -3 switch. The extra channel is the result of doing two interpolations of the Bayer green channel such that the greens on the same row as the reds produces one channel and the greens on the same row with the blues produce another channel. The second green channel will be the highest numbered channel in the output. This option is used to gauge the difference between greens in each row to decide whether the DNG BayerGreenSplit tag should be used for a given source of image data (e.g. camera).

The "-s <CFA index>" option chooses which set of color filter arrays to use when there are multiple ones for an input image. Each CFA array is a separate channel in the DNG input. This applies to the Fuji SR cameras for example, where the first channel is from the S sensing elements and the second channel is from the R sensing elements. The S elements are more sensitive and the R elements are less so with the goal of using both to increase the dynamic range the sensor can capture in a single image. By default dng_validate generates an image from only the S sensors. By using "-s 1" the R sensing elements' data can be used to construct the output image. (This index is 0 based. The default is 0.)

The "-q <target-binned-size>" option enables binning during the demosaic process. This is useful for creating previews or thumbnails. The binning factor is determined from the target-binned-size, which is the size in pixels of the larger dimension of the image that is desired. An integer binning factor will be computed to produce an image of that size or larger. For example, if the input image is 3008 x 2000 pixels and the target-binned-width is 700, factor of 4 binning will be used and the result output image (after demosaicing) will be 752 x 500 pixels.

The "-cs1" option generates the output image in sRGB color space.

The "-cs2" option generates the output image in AdobeRGB color space.

The "-cs3" option generates the output image in ProPhotoRGB color space.

The "-cs4" option generates the output image in ColorMatch color space.

The "-cs5" option generates the output image in grayscale gamma 1.8 color space.

The "-cs6" option generates the output image in a grayscale gamma 2.2 color space.

The "-16" option causes dng_validate to output 16-bit per component images rather than the default 8-bit.

The "-1" option causes the unprocessed raw image data to be written to the named output file. This applies only to the next input file after the switch. This switch can occur multiple times on a single command line.

The "-2" option causes the image data after linearization and black/white level mapping to be written to the named output file. This applies only to the next input file after the switch. This switch can occur multiple times on a single command line.

The "-3" option causes the image data after demosaic processing, but prior to color space conversion, noise reduction, sharpening, etc., to be written to the named output file. This applies only to the next input file after the switch. This switch can occur multiple times on a single command line.

The "-tif" option causes the final rendered image to be written as TIFF to the named output file. This applies only to the next input file after the switch. This switch can occur multiple times on a single command line.

The "-dng" option causes the parsed DNG data to be reserialized and written to the named output file. This mostly serves to provide an example code path for the process of writing a DNG file as the output may not differ significantly from the input DNG. (Parameters, such as whether the data is compressed or not, may vary between the input and output DNG files.) This applies only to the next input file after the switch. This switch can occur multiple times on a single command line.